

WHAT IS CLAIMED IS:

1. A method for simulating different mean time to recover (MTTR) settings,
the method comprising:

providing a simulated checkpoint queue, the simulated checkpoint queue
being associated with a simulated MTTR setting, the simulated
checkpoint queue being an ordered list of one or more elements,
each of the one or more elements representing a respective buffer,
the ordered list having a head and a tail;

in response to detecting a change to a first buffer in a normal checkpoint
queue, checking if the first buffer is represented in the simulated
checkpoint queue, and if the first buffer is not represented in the
simulated checkpoint queue, linking an element that represents the
first buffer to the tail of the simulated checkpoint queue.

2. The method of Claim 1 further comprising:

providing a simulated write counter, the simulated write counter being
associated with the simulated MTTR setting;

determining if linking the element to the tail of the simulated checkpoint
queue causes the simulated checkpoint queue to exceed a
predetermined length; and

in response to determining that the simulated checkpoint queue exceeds
the predetermined length, removing an element from the head of
the simulated checkpoint queue and incrementing the simulated
write counter.

3. The method of Claim 2, wherein the predetermined length being a dirty
buffer limit.

4. The method of Claim 1 further comprising, in response to detecting a write
out of a second buffer from volatile memory and storing in nonvolatile memory, checking
if the second buffer is represented in the simulated checkpoint queue, and if the second
buffer is represented in the simulated checkpoint queue, removing the element

representing the second buffer from the simulated checkpoint queue and incrementing the simulated write counter.

5. The method of Claim 4, wherein the write out of the second buffer being caused by an incremental checkpoint operation.

5 6. The method of Claim 1, wherein each element in the simulated checkpoint queue comprises:

- a first identifier that identifies an associated buffer; and
- a second identifier that identifies a journal entry in a redo log, the journal entry corresponding to the associated buffer.

10 7. The method of Claim 6, wherein the elements in the simulated checkpoint queue are ordered according to each element's journal entry position in the redo log.

8. The method of Claim 1 further comprising determining a dirty buffer limit for the simulated checkpoint queue, the dirty buffer limit specifying the length of the simulated checkpoint queue, the dirty buffer limit being determined from the simulated
15 MTTR setting and historical operating data.

9. The method of Claim 8, wherein the historical operating data comprises an average time to read one journal entry in a redo log.

10. The method of Claim 8, wherein the historical operating data comprises an average time to read one buffer from nonvolatile memory to volatile memory.

20 11. A computer-readable storage medium having stored thereon computer instructions that, when executed by a computer, cause the computer to:

- provide a simulated checkpoint queue, the simulated checkpoint queue being associated with a simulated mean time to recovery (MTTR) setting, the simulated checkpoint queue being an ordered list of one or more elements, each of the one or more elements representing a
25 respective buffer, the ordered list having a head and a tail;

in response to a change to a first buffer, check if the first buffer is represented in the simulated checkpoint queue, and if the first buffer is not represented in the simulated checkpoint queue, link an element that represents the first buffer to the tail of the simulated checkpoint queue.

12. The computer-readable storage medium of Claim 11 further storing computer instructions that, when executed by a computer, cause the computer to:

provide a simulated write counter, the simulated write counter being associated with the simulated MTTR setting; and

determine if linking the element to the tail of the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, remove an element from the head of the simulated checkpoint queue and increment the simulated write counter.

13. The computer-readable storage medium of Claim 12, wherein the predetermined length being a dirty buffer limit.

14. The computer-readable storage medium of Claim 11 further storing computer instructions that, when executed by a computer, cause the computer to, in response to a write out of a second buffer from volatile memory and storing in nonvolatile memory, check if the second buffer is represented in the simulated checkpoint queue, and if the second buffer is represented in the simulated checkpoint queue, remove the element representing the second buffer from the simulated checkpoint queue and increment the simulated write counter.

15. The computer-readable storage medium of Claim 14, wherein the write out of the second buffer being caused by an incremental checkpoint operation.

16. The computer-readable storage medium of Claim 10, wherein each element in the simulated checkpoint queue comprises:

a first identifier that identifies an associated buffer; and

a second identifier that identifies a journal entry in a redo log, the journal entry corresponding to the associated buffer.

17. The computer-readable storage medium of Claim 16, wherein the elements in the simulated checkpoint queue are ordered according to each element's journal entry position in the redo log.

18. The computer-readable storage medium of Claim 11 further storing computer instructions that, when executed by a computer, cause the computer to determine a dirty buffer limit for the simulated checkpoint queue, the dirty buffer limit specifying the length of the simulated checkpoint queue, the dirty buffer limit being determined from the simulated MTTR setting and historical operating data.

19. The computer-readable storage medium of Claim 18, wherein the historical operating data comprises an average time to read one journal entry in a redo log.

20. The computer-readable storage medium of Claim 18, wherein the historical operating data comprises an average time to read one buffer from nonvolatile memory to volatile memory.

21. A system comprising:

a memory;

one or more processors coupled to the memory;

a simulated MTTR setting maintained in the memory;

a simulated checkpoint queue maintained in the memory, the simulated checkpoint queue being associated with the simulated MTTR setting, the simulated checkpoint queue being an ordered list of one or more elements, each of the one or more elements representing a respective buffer, the ordered list having a head and a tail; and

a simulated write counter maintained in the memory, the simulated write counter being associated with the simulated MTTR setting, the simulated write counter providing a count of the number of times an element is removed from the simulated checkpoint queue, wherein the element is removed from the simulated checkpoint queue in response to a write out of a buffer from volatile memory and storing in nonvolatile memory.

22. The system of Claim 21, wherein the element removed from the simulated checkpoint queue represents the buffer written out from volatile memory and stored in nonvolatile memory.

23. The system of Claim 22, wherein the buffer is written out from volatile memory and stored in nonvolatile memory as a result of an incremental checkpoint operation involving the buffer.

24. The system of Claim 21, wherein the buffer is written out from volatile memory to nonvolatile memory as a result of a second element being linked to the simulated checkpoint queue, wherein linking the second element causes the simulated checkpoint queue to exceed a predetermined length.

25. The system of Claim 24, wherein the second element does not represent the buffer written out from volatile memory and stored in nonvolatile memory.

26. The system of Claim 24, wherein the second element is linked to the simulated checkpoint queue in response to a modification to a second buffer, the second buffer being represented by the second element.

27. The system of Claim 24, wherein the predetermined length being a dirty buffer limit, the dirty buffer limit being determined from the simulated MTTR setting and historical operating data.

28. The system of Claim 27, wherein the historical operating data comprises an average time to read one journal entry in a redo log.

29. The system of Claim 27, wherein the historical operating data comprises an average time to read one buffer from nonvolatile memory to volatile memory.

10017337 421201